# Amplify 2026

November 2-5 Scottsdale, AZ

# AI under the hood: Governance, architecture and defensible adoption



Wolters Kluwer

**ELM Solutions**

# Speakers



**Vincent Venturella**
Director,
Product Management
Wolters Kluwer



**Jaya Prasad Jayakumar**
Associate Director,
Data Science
Wolters Kluwer



**Bryan Rodriguez**
Technical
Product Manager,
MMC

# Takeaways

1. The basics of AI functionality
2. Governance and risk management
3. The essential questions to ask first

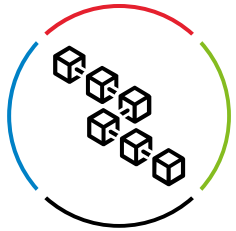




# AI basics

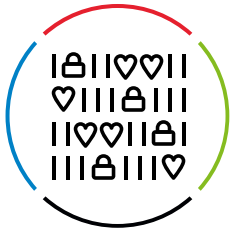
# 1 What an AI system actually is

A system combines models, data, instructions, tools, and people. The model is only one component.



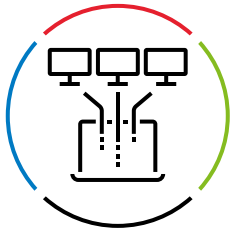
## Model

Recognizes patterns and produces likely outputs from the context provided



## Context + data

Grounds the model in enterprise information, permissions, and business meaning



## Tools + workflow

Allows the system to retrieve information or perform approved actions

**Human judgment remains the accountability layer.**

# 2 What AI can do, and what it cannot

Think “powerful assistant with variable judgment,” not “independent authority.”

## AI is strong at:

**Useful patterns**

- Summarize and draft
- Find and classify information
- Explain or compare options
- Surface patterns across content
- Trigger bounded, approved workflows

## AI is not reliably able to:

**Important limits**

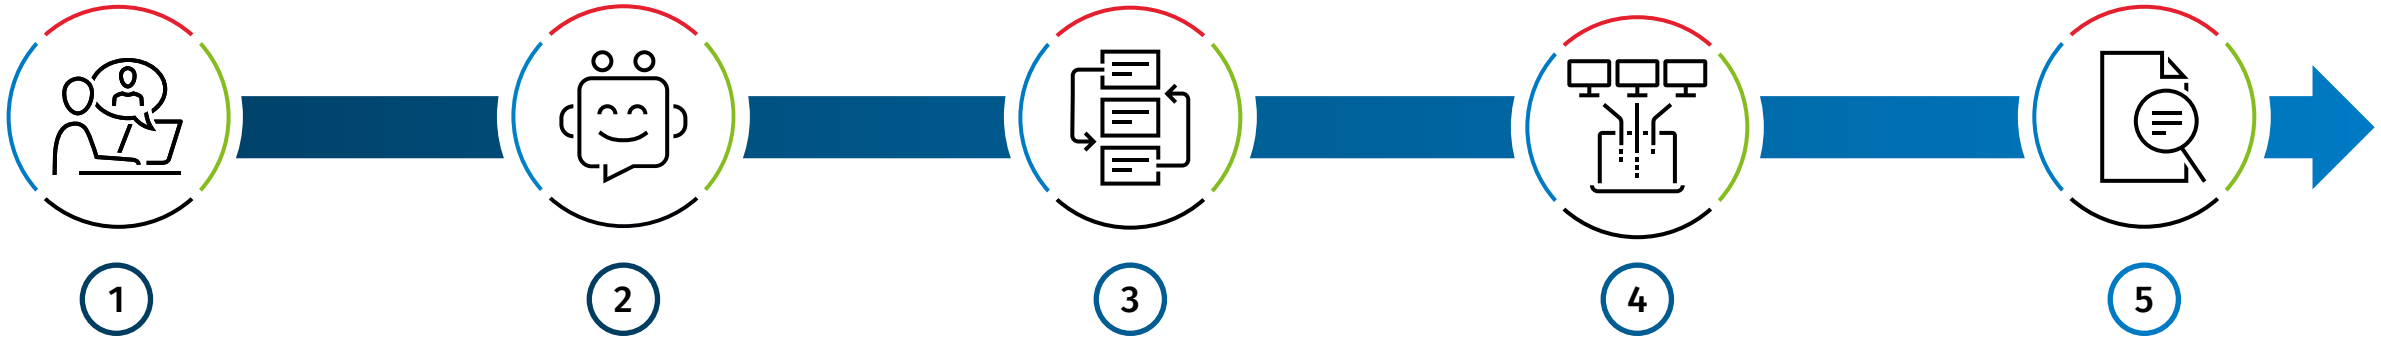
- Guarantee truth or completeness
- Understand unstated business nuance
- Resolve poor or conflicting data
- Own accountability for decisions
- Operate safely without controls



**Best use** Accelerate people, preserve professional review for consequential outcomes.

# 3 Agents and MCP servers in plain language

Agents organize work. MCP provides a standard doorway to approved tools and data.



1. **User goal**
   "Help me prepare for this matter"
2. **Agent**
   Plans steps and chooses tools
3. **MCP gateway**
   Applies identity, permissions, logging
4. **Enterprise tools**
   Retrieves data or performs bounded actions
5. **Reviewed outcome**
   Human verifies and approves

**Business takeaway**

An MCP server does not make data public. Its value is a consistent, governable connection point where identity, access, approved tools, and monitoring can be enforced.



# The data challenge

# 4 The operating model: Three foundations, one outcome

AI scales when data, governance, and risk management reinforce each other.



1

## Aligned data

Trusted, connected, permission-aware context

2

## Governance framework

Clear ownership, stage gates, and lifecycle controls

3

## Risk management

Risk-based safeguards, evidence, and monitoring

**Scalable • Secure • Defensible AI**

# 5 Why disconnected data becomes an AI problem

AI amplifies the context it receives. Fragmented context creates fragmented answers.




### Different definitions

"Matter," "customer," or "approved" may mean different things across systems.

### Incomplete context

The answer may omit records that live in another repository or format.

### Inconsistent permissions

A technically connected source may still expose more than the user should see.

The goal is not "put everything in one database." The goal is governed, understandable access across sources.

# 6 Create an AI-ready context layer

Connect meaning, identity, permissions, and lineage without forcing a single-system rewrite.

| 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|
| Business meaning | Identity and access | Quality and provenance | Integration | Observability |
| Common definitions, owners, and approved use | The AI inherits the user's entitlements | Know where information came from and when it changed | APIs, retrieval, and MCP expose approved capabilities | Log requests, sources, tool calls, and outcomes |

**Design principle** Expose the minimum data and actions needed for the use case.



# Risk governance

# 7 Governance turns an experiment into an enterprise capability

Governance is the decision system around AI, not paperwork added at the end.

| | |
|---|---|
| **Intake** | What outcome, user, and data? |
| **Assess** | Risk, value, feasibility, vendor |
| **Approve** | Named owner and stage-gate decision |
| **Build + test** | Controls, evaluation, documentation |
| **Deploy + monitor** | Usage, quality, incidents, change |

**Evidence travels with the system across its lifecycle.**

# How I set controls before delegating

**Defined procedure**
Programmed rules

**Adapt and interpret**
AI



**What did I allow it to do?**



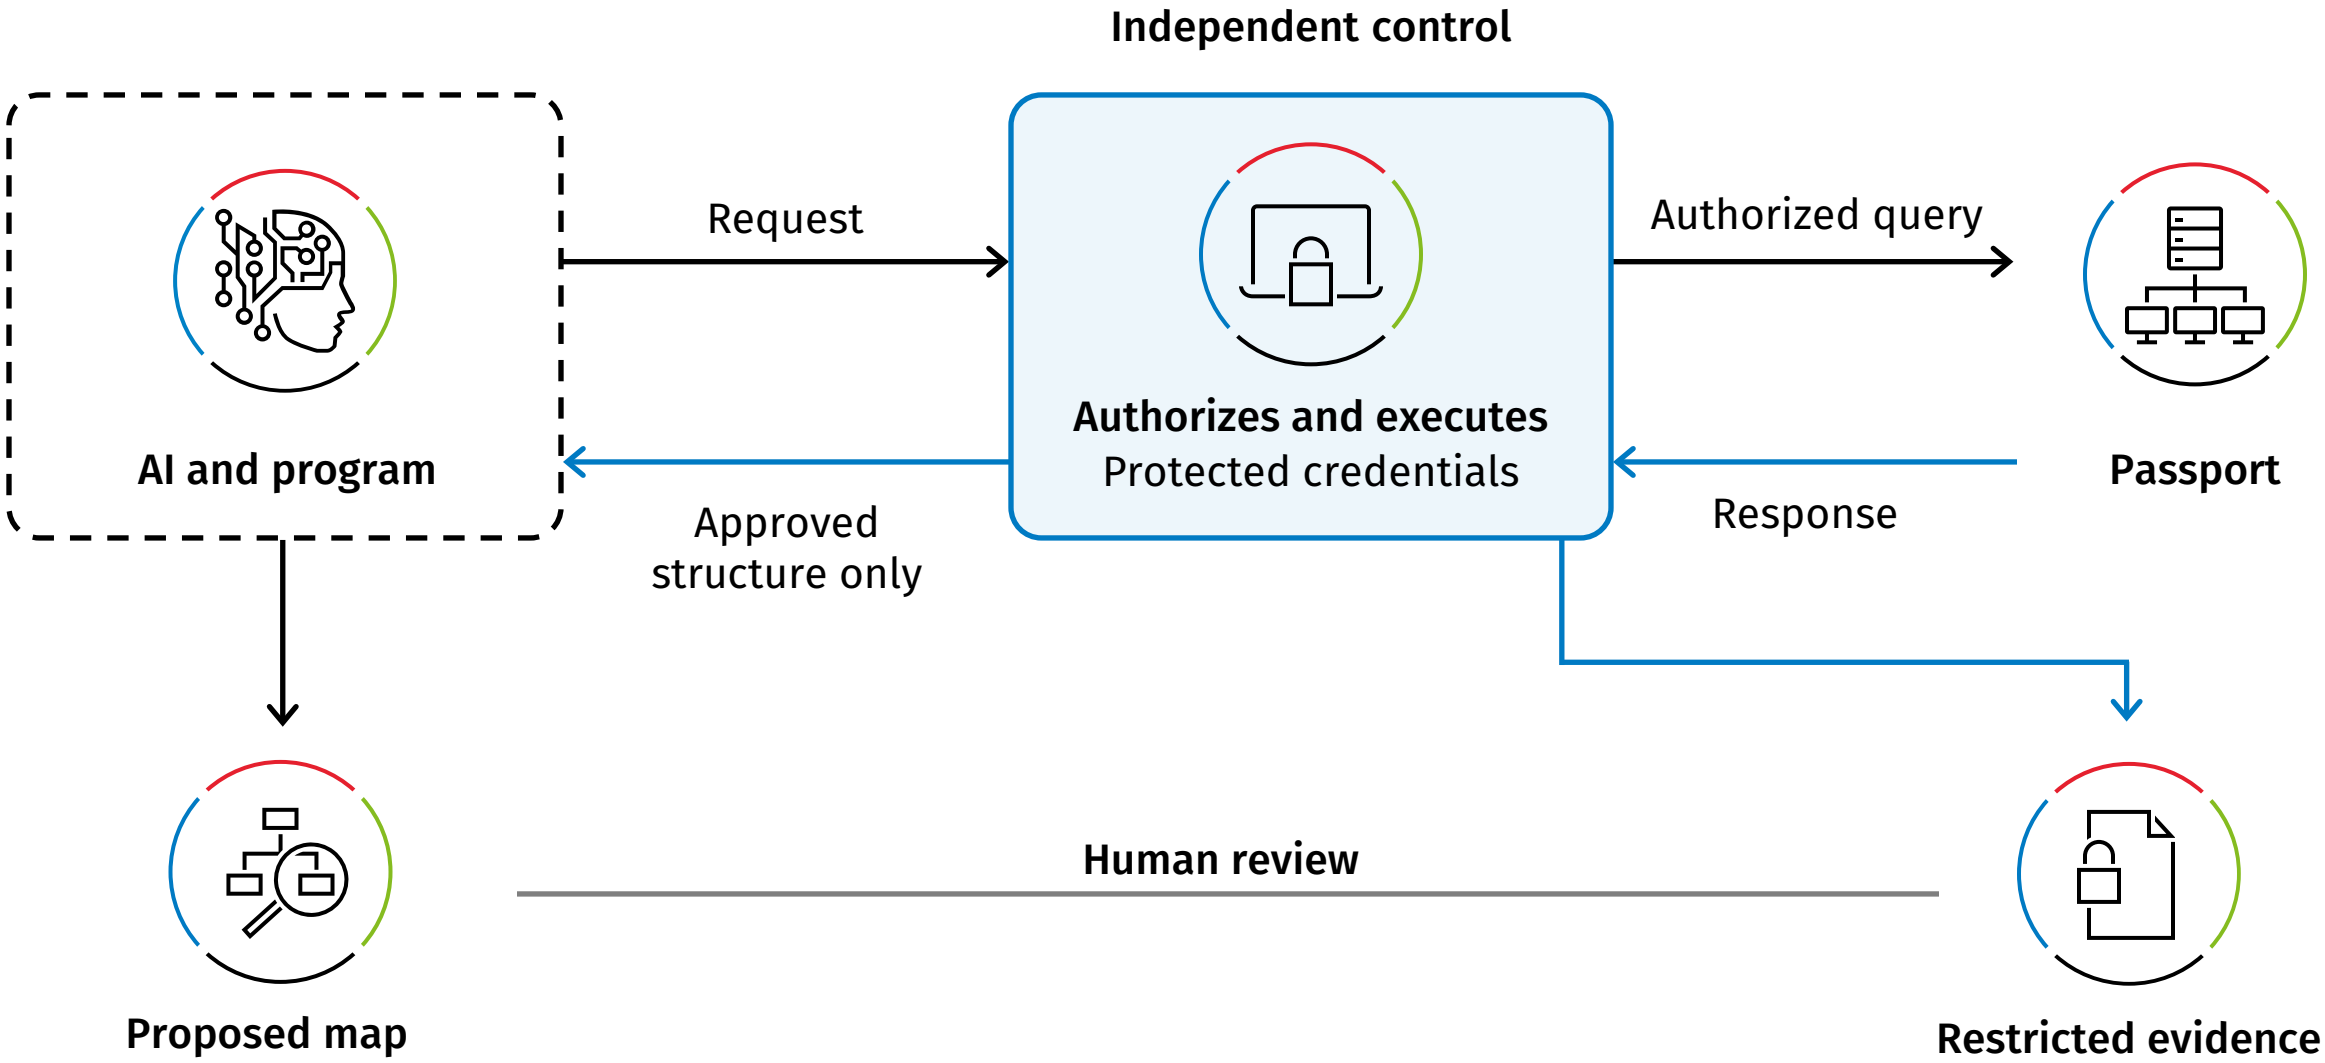
**What information did I share?**



**What did I require to accept the result?**

**Controls before delegation**

# Passport: How I kept control



# 8 Would you let the AI act on this?

Illustrative invoice review: An expense requires prior approval.

## What the AI can verify

The expense appears on the invoice.

No approval appears in the available records.

What should happen next?

A Deduct the expense

B Treat it as a confirmed violation

C Request evidence and hold the decision

**An unsuccessful search does not prove approval never existed.**

# 9 Evaluation against real business decisions

Domain experts define the expected response, including when the AI should ask for help.

| Evaluation case | Expected system response |
| --- | --- |
| Clear violation with evidence | Flag the issue and cite the supporting evidence. |
| Documented, approved exception | Recognize the exception and avoid an incorrect flag. |
| Approval evidence is missing | Explain what is unverified and request the missing evidence. |
| Applicable policies conflict | Surface the conflict and route the decision for resolution. |



**Measures that reveal decision quality**

Missed issues, incorrect flags, appropriate escalation, and reviewer effort.

# 10 Readiness is specific to the task

Each responsibility requires evidence that the system can perform it reliably.

1

## Summarize

- Evidence for an accurate summary
- Faithful summaries
- Supported claims
- Visible information gaps

2

## Recommend

- Additional evidence for a proposal
- Correct policy application
- Exceptions recognized
- Acceptable error rates

3

## Execute

- Additional evidence for an action
- Tested action controls
- Verified execution
- Tested recovery

**What evidence would earn the next level of responsibility?**

# 11 Risk management begins with a shared risk language

Risk is not only model accuracy. It spans the full sociotechnical system.

**Information**
Privacy, confidentiality, access

**Reliability**
Hallucination, drift, incomplete context

**Security**
Prompt injection, tool abuse, data leakage

**Legal + compliance**
Obligations, records, transparency

**Human + process**
Over-reliance, unclear accountability

**Operational**
Resilience, cost, vendor dependency

**Assess the whole workflow: Input → model → tools → action → human outcome.**

# 12 A pragmatic path from use case to scale

A system combines models, data, instructions, tools, and people. The model is only one component.

1 2 3 4 5



## Choose

- Concrete workflow
- Named owner
- Measurable outcome



## Prepare

- Map data + permissions
- Assign risk tier
- Define “good”



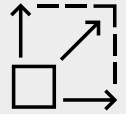
## Pilot

- Small user group
- Human review
- Capture failures



## Prove

- Evaluate quality
- Test controls
- Document decisions



## Scale

- Monitor continuously
- Manage change
- Expand deliberately

**Increase autonomy only when evidence shows the system and its controls perform as intended.**

# The decision to make next

Choose one use case and answer these five questions before discussing models or vendors.

1. What business outcome matters?
2. Which data and actions are truly required?
3. Who owns the system and the decision?
4. What could go wrong, and how will we know?
5. **What evidence is required before scaling?**



**SUCCESS FORMULA**

Aligned data
+
Governance
+
Risk management

**Scalable • Secure • Defensible**

# Thank you!



Wolters Kluwer